

The shifted proper orthogonal decomposition A mode decomposition for multiple transport phenomena

Julius Reiss* Philipp Schulze† Jörn Sesterhenn*

December 8, 2015

Abstract

Transport-dominated phenomena provide a challenge for common mode-based model reduction approaches. We present a model reduction method, which extends the proper orthogonal decomposition (POD) by introducing time-dependent shifts of the snapshot matrix. The approach, called shifted proper orthogonal decomposition (sPOD), features a determination of the *multiple* transport velocities and a separation of these. One- and two-dimensional test examples reveal the good performance of the sPOD for transport-dominated phenomena and its superiority in comparison to the POD.

1 Introduction

Model reduction is a vital tool to derive low-dimensional models from experimental or numerical data. These reduced-order models provide fast simulation tools commonly used for control, optimization, and parameter studies and are nowadays a necessary tool in many fields of applications, as in electrical engineering, fluid dynamics, and structural dynamics, etc.

Most of the model reduction methods for linear or nonlinear systems are based on the idea of projecting the high dimensional solution on a low-dimensional subspace, which is spanned by appropriate ansatz functions or modes. The number of modes needed to obtain a sufficiently accurate approximation of the system in question decides on the computational cost and thereby on the usefulness of the reduced system. Transport-dominated phenomena are usually a challenge for these mode-based methods, since their dynamical behavior cannot be captured accurately by a few spatial modes (cf. [6, 19]).

The contribution of this paper is the introduction of a model reduction approach, which is able to provide low-dimensional representations of transport-dominated phenomena. In addition, it is a generalization of the proper orthogonal decomposition (POD) and so it performs equally good when considering model reduction tasks, where common approaches as the POD are used with success. The key tool is a time-dependent shift in combination with a procedure to separate different transports within the system. The dominant transport velocities are determined by considering the dependence of the singular values of the shifted snapshot matrix on the time-dependent shift.

*Institut für Strömungsmechanik und Technische Akustik, TU Berlin, Germany, {reiss,sesterhenn}@tnt.tu-berlin.de

†Institut für Mathematik, TU Berlin, Germany, pschulze@math.tu-berlin.de.

It should be mentioned that the idea of shifting the solution is also used in the framework of symmetry reduction (cf. [4, 10, 23]). However, in contrast to that framework, the method presented here is purely based on data and does not assume self-similarity of the solution. Moreover, the separation of two or more transports has not been addressed up to now.

We focus on obtaining a low-dimensional representation of the solution. The corresponding low-dimensional subspace can be used to build a reduced order model (ROM), for instance by an interpolation procedure as in [17] or by a Galerkin projection of the original model. Both approaches will be subject of upcoming research.

The possibility to save computation time and/or memory in design and control applications, and the chance to obtain a deeper understanding of the physical mechanisms has led to an extensive research in the area of reduced-order models in the past decades. In general, the goal is to obtain a low-dimensional model, which approximates the mapping from a set of inputs to a set of outputs. Among others, inputs can be design parameters, system conditions, or controls. Common outputs are, for instance, performance or physical quantities like measurements or even full flow solutions. A major class of reduced-order modeling approaches is given by input-output interpolation methods, which do not aim in reducing the internal dynamics but which are only based on the input-output behavior, usually described by the transfer function (see [1, 3, 14, 20] and references therein). These methods are successfully applied to linear systems, but extensions to nonlinear systems are rare and have some drawbacks as low computational efficiency (cf. [11]). In contrast, the most common model reduction techniques for nonlinear systems are based on modes. Examples are reduced basis methods [8, 9, 13, 24, 28], dynamic mode decomposition [5, 25, 26], and the POD [15, 21, 29]. All these methods have in common that the state variables are described by a superposition of modes.

The most popular method for mode construction from data is the POD, where the modes are determined from the singular value decomposition (SVD) of a snapshot matrix. The SVD provides the best rank r approximation (where r refers to the chosen number of modes) in the sense that it minimizes the error in the 2-norm. This will be briefly discussed in section 2. Depending on the application, the POD has several synonyms, as for instance, principal component analysis [16] or Karhunen-Loeve decomposition [7].

The 2-norm optimality obtained by the POD should not be mistaken for the optimal decomposition in general. If the input-output dynamics is mediated by a small subspace of the dynamical system, its contribution to the norm might be negligibly small, yet be of major interest. This aspect is accounted for at balanced model reduction techniques as the balanced proper orthogonal decomposition [27, 29]. This method also considers the adjoint dynamics, which describes the influence of a modification. Thus, the modes obtained by this method capture the input-output behavior of the system, but do not necessarily give a good approximation of the internal states. The balanced POD is restricted to linear systems.

Interestingly, the POD does not include causality. Reordering the snapshots leads to exactly the same spatial modes. Further, while orthogonality is attractive for many reasons, it might be misleading when the operator of the system is non-normal. Non-normality of the eigenmodes is an important aspect of many fluid dynamical systems. The dynamic mode decomposition (DMD) aims at these two aspects by calculating an approximate linear operator, which describes the dynamics of a given system. The obtained modes might or might not be orthogonal and, furthermore, include causality.

While the different mode-based methods address different aspects of the system, transport-dominated phenomena represent a challenge for all of them. Even for simple transport phe-

nomena, as a shock wave with constant velocity, the common mode-based methods fail in obtaining an accurate and low-dimensional approximation. The main problem is that the mode decompositions are usually formulated as a linear combination of spatial modes with time-variant coefficients. However, this structure is not suitable for transport-dominated phenomena.

This motivates research for inventing a more general mode decomposition. Two principles guide our effort. First, we require that the new method generalizes its base method, the POD, so that the low-dimensional approximation is at least as efficient as when applying the common POD. Second, we want to build a purely data-based method, which does not depend on any additional structure like hyperbolicity.

In the following section, we motivate and derive the method by means of a simple example in a one-dimensional domain, which still is a challenge for common model reduction techniques as POD. For this, first we consider one moving signal and show how to reduce it. Then, a more complex system with two different transports is considered and a procedure is derived to separate the different velocity components. Then the detection of the velocities is discussed. In section 3 we apply the new method to a two-dimensional test case from fluid mechanics with transported developing vortex pairs with non-trivial velocities. Finally, we conclude and give an outlook on ongoing and planned work.

2 One dimension

In this section the method will be developed considering a one-dimensional example. As a test problem, the linear wave equation

$$\begin{aligned}\partial_t \rho + \rho_0 \partial_x u &= 0, \\ \partial_t u + c^2 / \rho_0 \partial_x \rho &= 0,\end{aligned}\tag{1}$$

is considered on a periodic domain $x = (0, L]$. Here, u is the velocity, ρ is the density (fluctuation), ρ_0 is the reference density and c the speed of sound. For the remainder of this paper, we set $\rho_0 = 1$ and $c = 1$. The general solution can be written as

$$q(x, t) = \begin{bmatrix} \rho(x, t) \\ u(x, t) \end{bmatrix} = q_+(x - ct) \begin{bmatrix} \rho_0 \\ c \end{bmatrix} + q_-(x - (-c)t) \begin{bmatrix} \rho_0 \\ -c \end{bmatrix}\tag{2}$$

with arbitrary initial conditions $q_{\pm}(x)$ for the two transported quantities, the Riemann invariants. The corresponding transport velocities are $c_{\pm} = \pm c$. By choosing q_{\pm} proportional to $\exp(2n\pi i(x \pm ct)/L)$, the solution can be rewritten as

$$q(x, t) = \sum_{n=1}^{\infty} \left(\alpha_n \sin(2\pi n x) \begin{bmatrix} \rho_0 \cos(2\pi n c t + \phi_n) \\ c \sin(2\pi n c t + \phi_n) \end{bmatrix} + \beta_n \cos(2\pi n x) \begin{bmatrix} \rho_0 \cos(2\pi n c t + \gamma_n) \\ c \sin(2\pi n c t + \gamma_n) \end{bmatrix} \right)\tag{3}$$

with constants α_n , β_n , ϕ_n , and γ_n . Every non-zero amplitude α_n, β_n yields a vibrating mode of a string.

As mentioned in section 1, model reduction methods often build on describing a dynamical system by a superposition of a few modes. One of the most popular approaches is the POD,

which aims at representing the solution as a linear combination of orthonormal modes ϕ_k

$$q(x_i, t_n) = \sum_k \alpha_k(t_n) \phi_k(x_i). \quad (4)$$

with time-dependent coefficients α_k . Usually, the snapshot matrix $X_{ij} = q(x_i, t_j)$ is introduced, which is a space-time-discretized solution. A low-dimensional representation minimizing the approximation error in the 2-norm is demanded, which is given by the singular value decomposition (SVD) of the snapshot matrix X . Often usage of the full snapshot matrix is omitted to circumvent the enormous numerical effort [2]. To this end, a reduced set of discrete time and spatial points is used. This does not effect the quality of the obtained modes significantly as long as the reduced set of discrete points still captures the dynamical behavior of the system. This reduced snapshot matrix is often much smaller than the resolution needed for a reliable simulation. The singular values determine the mean square amplitude of the corresponding modes, which coincides with the kinetic energy if the mode represents the velocity [18]. Consequently, strongly decaying singular values allow a good representation with a few modes.

The ability of the POD to describe the solution of (1) strongly depends on the initial condition. If a vibrating string with a dominant frequency and a few harmonics is chosen, i.e., as in (3) with a few dominant α_n, β_n , the POD will find these few modes and, consequently, it will deliver an accurate low-dimensional representation. However, if a transported quantity with steep gradient is given, i.e., as in (2) with strongly localized initial conditions q_{\pm} , the singular value decay is rather gradual and one is forced to use a high number of modes to get a reasonable description. Thus, by setting the initial conditions, the POD approximation can be changed from exemplary to failing. This variability makes the example well-suited to develop a method, which can handle transported quantities while including the classical POD as a special case.

2.1 One transported quantity

Setting $q_-(x) = 0$ in the solution (2) of the linear wave equation (1) describes a single transported quantity, which is given by a shift of the initial condition q_+ with the transport velocity c . To provide a challenging case for POD, a sharp Gaussian pulse of $q_+(x) = \exp(-(x - x_0)^2/\delta^2)$ is chosen with $\delta = L/50$. The analytic solution is shown in Fig. 1 (left) for the time interval $[0, 1.25L/c]$. For the discrete snapshot matrix $X_{ij} = q(x_i, t_j)$, we chose 250 equidistant time and 200 space points. Often the snapshot matrix contains only a subset of the system variables, either by choice or, if based on experimental data, by availability. In this subsection we only use the density, if not stated otherwise.

For the considered example, a very slow decay of the singular values is found (cf. Fig. 1, left), resulting in the need of a high number of modes to represent the solution. Even though the analytic solution can be formulated by only one transported Gaussian pulse, many POD modes are needed. The reason for this is that the structure (4) is not adequate to describe a solution of the form (2). Roughly speaking the dyadic products describe rectangular (with respect to the space time diagram) structures well, while diagonal structures such as in Fig. 1, left, cannot be easily written as a dyadic product. Transport-dominated phenomena are reflected by diagonal structures in the space time diagram and need to be approximated by the combination of many rectangular structures of different size. This illustrates the problem of approximating these phenomena by a small number of time-invariant spatial modes. However,

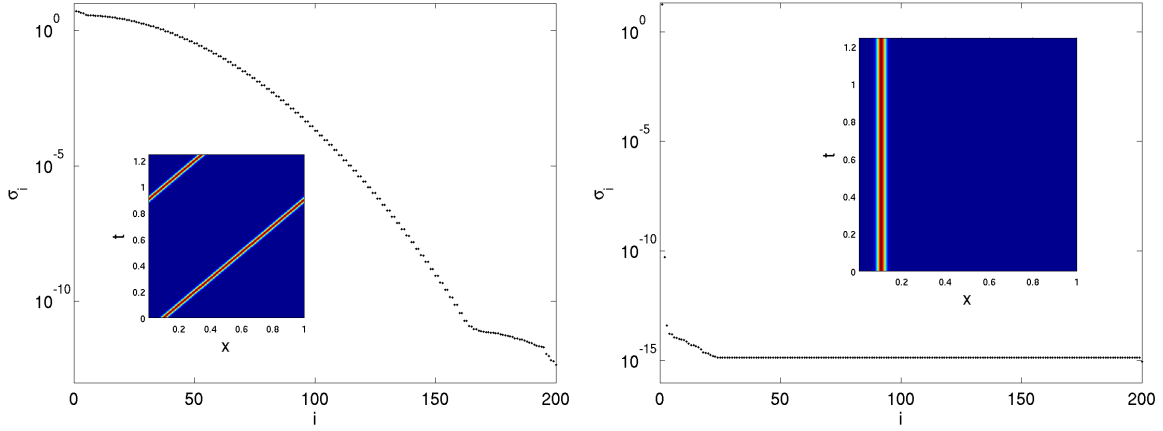


Figure 1: Single transported quantity. Left: Density in the space time diagram. The spectrum shows a very slow decay. Right: If the solution is described in a co-moving frame, i.e., the solution is shifted at every time to compensate the transport, the singular spectrum shows a rapid decay. One mode is sufficient to describe the flow (apart from a small numerical error of the shift procedure).

if the velocity is known or can easily be detected, the solution is near at hand. A time-dependent shift which compensates the transport velocity $T^{-c}(q(x, t)) = q(x, 0)$ yields a structure that can be written as a single dyadic product, which in turn corresponds to the description by just one mode, Fig. 1, right. In a similar manner, shifts have been used in other reduction frameworks, e.g., for the model reduction of a combustion [17] and in a symmetry reduction framework [4, 10, 23].

2.2 Multiple transported quantities

Many systems from real-world applications, as aerodynamics or chemical engineering, do not have just one, but multiple transported quantities. To approach these more challenging problems, we start by considering the solution (2) of the wave equation (1) with

$$q_+ = q_- = \exp(-(x - L/2)^2/(L/50)^2),$$

which initially describes a pressure pulse in a system at rest. The analytic solution for these two transported quantities is shown in Fig. 2. Due to the solution structure (cf. (2)), it is known, that the full information is described by just two modes. Ideally, a model reduction approach should detect these modes. However, while the solution structure is known in this special case, we refrain from using details of it, since we are looking for a general procedure. The desired method should find such modes purely data-based. In this way, such a procedure has the potential to work for data of non-hyperbolic transport, too. Nevertheless, for now we assume the transport velocities to be known and rather focus on detecting and separating the two transported quantities from the data. The determination of unknown transport velocities is discussed in section 2.3.

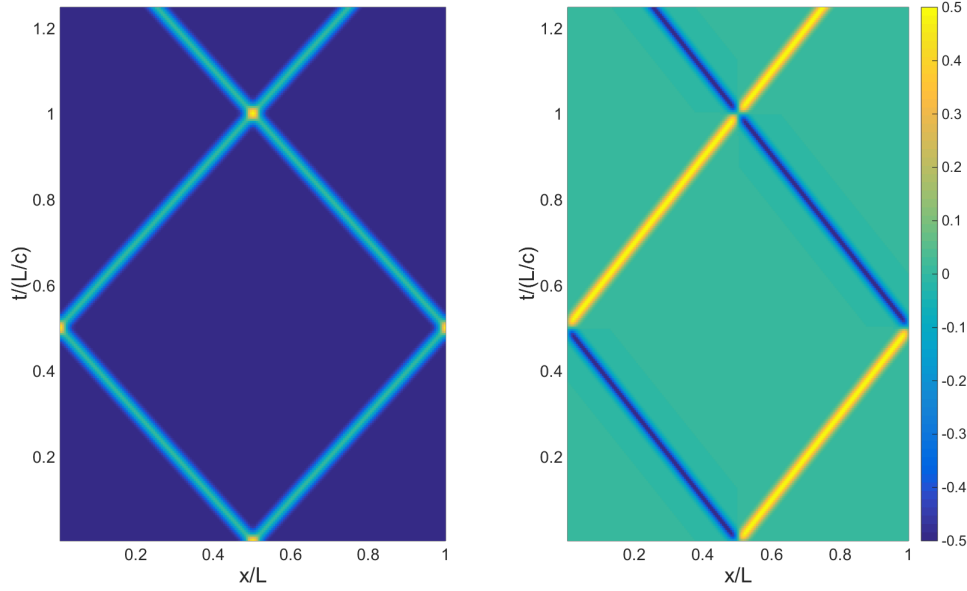


Figure 2: Two transported quantities: solution in the space time diagram (left: density, right: velocity)

In subsection 2.1 it was found that the transported phenomena are well described by the first modes in a co-moving system, while a decomposition of structures in a different velocity frame leads to the need of many modes to describe the dynamics reasonably well. Conversely, the first mode contains comparably little of the dynamics, if the shift velocity does not agree with one of the transport velocities. Thus, a naive approach to *decompose* different velocity components is given by the following procedure. For the sake of simplicity it is stated for two transport velocities but the extension to more velocities is straightforward.

THE NAIVE METHOD

Input:

- $k \times n$ snapshot matrix X with n (temporal) snapshots of k grid points
- transport velocities a, b with $a \neq b$

Procedure:

1. Shift & Decompose

Compute the SVDs

$$T^{-a}(X) = U_{X,a} \Sigma_{X,a} V_{X,a}^T$$

$$T^{-b}(X) = U_{X,b} \Sigma_{X,b} V_{X,b}^T$$

with $U_{X,a}, U_{X,b} \in \mathbb{R}^{k,k}$, $\Sigma_{X,a}, \Sigma_{X,b} \in \mathbb{R}^{k,n}$, and $V_{X,a}, V_{X,b} \in \mathbb{R}^{n,n}$. T^{-v} denotes the shift operator, which performs a shift of a snapshot matrix such that the shifted matrix corresponds to the solution in the co-moving reference frame with velocity v .

2. Truncate

Approximate $U_{X,i} \Sigma_{X,i} V_{X,i}^T$ by neglecting the singular values $\sigma_{r+1}, \dots, \sigma_{\min(k,n)}$ where r is chosen as small as possible, as great as necessary to obtain a good approximation. This leads to the approximate SVDs

$$\begin{aligned} T^{-a}(X) &\approx \tilde{U}_{X,a} \tilde{\Sigma}_{X,a} \tilde{V}_{X,a}^T \\ T^{-b}(X) &\approx \tilde{U}_{X,b} \tilde{\Sigma}_{X,b} \tilde{V}_{X,b}^T \end{aligned}$$

with $\tilde{U}_{X,a}, \tilde{U}_{X,b} \in \mathbb{R}^{k,r}$, $\tilde{\Sigma}_{X,a}, \tilde{\Sigma}_{X,b} \in \mathbb{R}^{r,r}$, and $\tilde{V}_{X,a}, \tilde{V}_{X,b} \in \mathbb{R}^{n,r}$.

3. Shift Back & Combine

Low-dimensional approximation

$$X \approx \tilde{X} = T^a \left(\tilde{U}_{X,a} \tilde{\Sigma}_{X,a} \tilde{V}_{X,a}^T \right) + T^b \left(\tilde{U}_{X,b} \tilde{\Sigma}_{X,b} \tilde{V}_{X,b}^T \right)$$

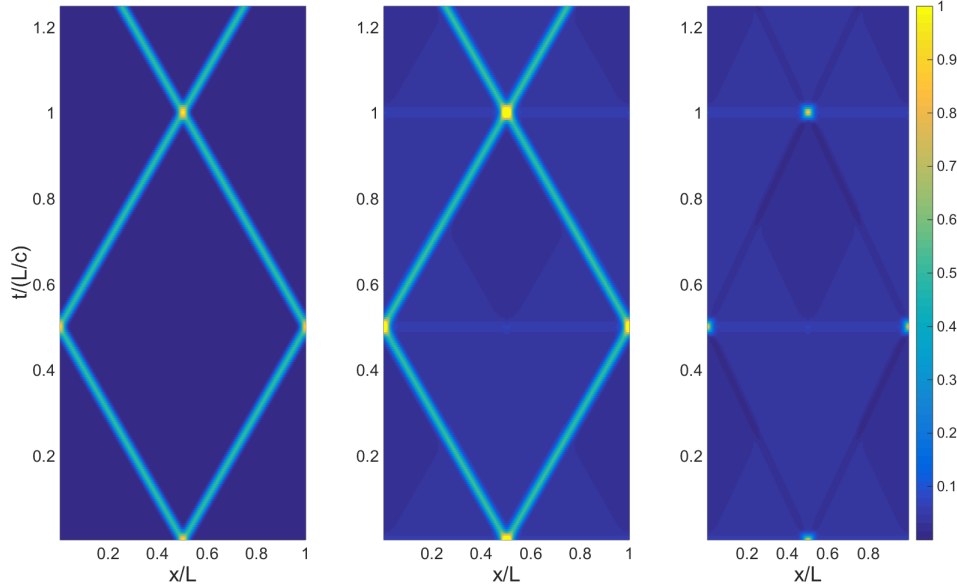


Figure 3: The density of two transported quantities, the left and right going wave: The comparison of full-order solution with the naive decomposition approximation (left: full-order solution, middle: naive decomposition approximation (2 modes), right: error)

The combination of the steps 1 and 2 leads to an approximate identification of the respective velocity component, since this should dominate the first modes, while the other velocity component gives a relatively small contribution. In step 3 the components are shifted back to the

original zero-velocity frame and combined. The numerical effort may be reduced by replacing the SVDs in step 1 with truncated SVDs [12]. We apply the naive approach to the density snapshot matrix depicted on the left-hand side of Fig. 2. The resulting low-dimensional approximation obtained by the naive approach is shown in Fig. 3.

While the main structures are reproduced well, the result is far from perfect. Especially, in the region where the left and the right going pulse overlap, a large error is produced. These strongly localized structures with respect to space and time are represented well in both velocity frames. After all, a point structure remains a point structure for any frame shift and can be described perfectly by a dyadic product. This ambiguity counteracts the separation of the two transported quantities when using the naive method. Furthermore, the quality of the solution cannot be improved by adding modes, since this countervails the separation.

To improve the naive algorithm, we introduce a mechanism of cleaning each transported quantity from the spurious influence of the other. To that end, for each transport we create a small ansatz space, which likely contains modes to build a proper decomposition. Precisely, we decompose the error, i.e., the difference between the original (X) and the approximated snapshot matrix (\tilde{X}), by the same naive method. The idea is that the double accounted quantities, which we want to remove from our representation, are contained in these error modes, since they arise from structures, which are represented well in all velocity frames. Thus, it remains to determine the best approximation by the naive modes and the error modes which reduces the problem to an optimization. This leads to the main algorithm:

THE SPOD ALGORITHM

1.-3. Create a naive decomposition as before

4. Calculate Error Modes

Calculate the error $E = X - \tilde{X}$, and from this compute the SVDs

$$\begin{aligned} T^{-a}(E) &= U_{E,a} \Sigma_{E,a} V_{E,a}^T \\ T^{-b}(E) &= U_{E,b} \Sigma_{E,b} V_{E,b}^T \end{aligned}$$

with $U_{E,a}, U_{E,b} \in \mathbb{R}^{k,k}$, $\Sigma_{E,a}, \Sigma_{E,b} \in \mathbb{R}^{k,n}$, and $V_{E,a}, V_{E,b} \in \mathbb{R}^{n,n}$.

5. Assemble Ansatz Basis

From the modes obtained in step 1 and the error modes obtained in step 4, set up the ansatz basis

$$B_a = \left\{ [U_{X,a}]_i ([V_{X,a}]_i)^T, [U_{E,a}]_i ([V_{E,a}]_i)^T \right\},$$

with $i = 1, \dots, s$. $[U_{X,a}]_i$ and $[U_{E,a}]_i$ denote the i th columns of $U_{X,a}$ and $U_{E,a}$, respectively, and so on. The number s is a chosen small number of modes comparable to r from step 2. The number of basis functions is $2s$. The ansatz basis for the second transport velocity, B_b , is constructed analogously.

6. Minimize Error Functional

Note that each of the ansatz bases B_a and B_b consists of $2s$ basis functions, which are matrices of dimension $k \times n$. These ansatz functions are denoted with $\Phi_{a,i}$ and $\Phi_{b,i}$, respectively, with $i = 1, \dots, 2s$. Now, construct the linear combination of the back-shifted basis functions that minimizes the error, i.e., solve the minimization problem

$$\min_{\alpha_{a,i}, \alpha_{b,i}} J(\alpha_{a,i}, \alpha_{b,i}) = \left\| X - \sum_{i=1}^{2s} \alpha_{a,i} T^a(\Phi_{a,i}) - \sum_{i=1}^{2s} \alpha_{b,i} T^b(\Phi_{b,i}) \right\|_2.$$

This minimization problem can be solved by reshaping and subsequent solving of an overdetermined linear equation system in a least squares sense.

7. Reduce Every Velocity Component

Calculate the SVDs

$$\begin{aligned} \sum_{i=1}^{2s} \alpha_{a,i} \Phi_{a,i} &= U_{\Phi,a} \Sigma_{\Phi,a} V_{\Phi,a}^T \\ \sum_{i=1}^{2s} \alpha_{b,i} \Phi_{b,i} &= U_{\Phi,b} \Sigma_{\Phi,b} V_{\Phi,b}^T \end{aligned}$$

and truncate them as before to obtain the approximations

$$\begin{aligned} \sum_{i=1}^{2s} \alpha_{a,i} \Phi_{a,i} &\approx \tilde{U}_{\Phi,a} \tilde{\Sigma}_{\Phi,a} \tilde{V}_{\Phi,a}^T, \\ \sum_{i=1}^{2s} \alpha_{b,i} \Phi_{b,i} &\approx \tilde{U}_{\Phi,b} \tilde{\Sigma}_{\Phi,b} \tilde{V}_{\Phi,b}^T. \end{aligned}$$

8. Low-Dimensional Representation

Eventually, the low-dimensional representation is obtained by assembling

$$X \approx \tilde{X} = T^a(\tilde{U}_{\Phi,a} \tilde{\Sigma}_{\Phi,a} \tilde{V}_{\Phi,a}^T) + T^b(\tilde{U}_{\Phi,b} \tilde{\Sigma}_{\Phi,b} \tilde{V}_{\Phi,b}^T).$$

The idea of step 4 is that the error modes should contain the structures, which are resolved by both velocities, and, thus, likely allow a linear combination to remove this doubled representation. Again, the numerical effort may be reduced by replacing the SVDs by truncated SVDs. In step 5 the ansatz bases are assembled. To improve the separation of the two transported quantities, the steps 4 to 8 are repeated iteratively. At the end of each iteration the approximate snapshot matrix \tilde{X} is updated and the modes U_X are replaced by the modes U_Φ computed in step 7. Note that even though the algorithm is formulated for two transports, the general procedure does not require this. The convergence behavior can be seen from the blue graph of Fig. 4 where the convergence is indicated by the decreasing mean error over

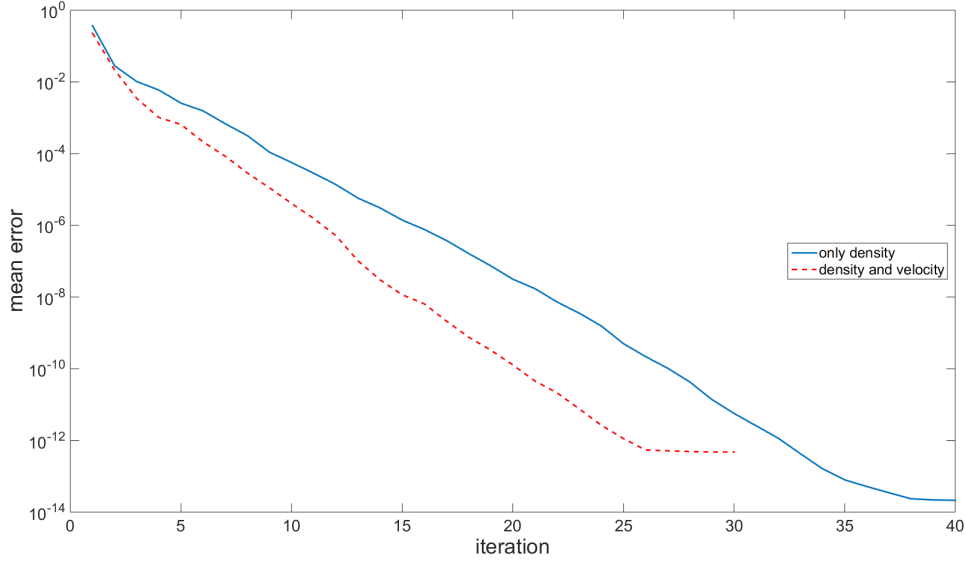


Figure 4: Two transported quantities: convergence of the sPOD algorithm (2 modes). The error of the sPOD approximation decreases to less than 10^{-12} and is most likely caused by the shift, which is accompanied by a numerical error due to an interpolation to grid points.

iterations. The considered error measure is given by the 2-norm of the error divided by the 2-norm of the full order snapshot matrix, i. e.,

$$\text{mean error} = \frac{\|X - \tilde{X}\|_2}{\|X\|_2}.$$

At that point it should be mentioned, that cross terms $U_i V_j^T$ with $i \neq j$ could be added. They often lead to a faster convergence of the separation. However, they are ignored here for reasons of computational effort. If the cross terms would be added, the resulting system of linear equations to be solved in step 6 would be of dimension $2kns^2$ instead of $2kns$. This is a minor issue for the one-dimensional example but it leads to high memory requirements and computation times when dealing with the two-dimensional example presented in section 3. In respect of examples of higher dimension, this bottleneck becomes even more inhibiting. For that reason the general algorithm is formulated without cross terms. Finally, the velocity subspaces are reduced and combined in steps 7 and 8, respectively, to obtain a low-dimensional representation.

The approximation obtained by the sPOD algorithm after 40 iterations is shown in Fig. 5. As can be seen, the solution is approximated excellently by just using two modes. This is the desired low-dimensional representation we have been looking for, since we know from the analytic solution that it is decomposed into two modes.

The good performance of the sPOD algorithm gets more striking when comparing it to the standard POD. For this purpose, Fig. 6 depicts the comparison between the full-order solution, the sPOD approximation with two modes, and the POD approximation with two modes. It is obvious that the POD approximation is highly inadequate, whereas the sPOD

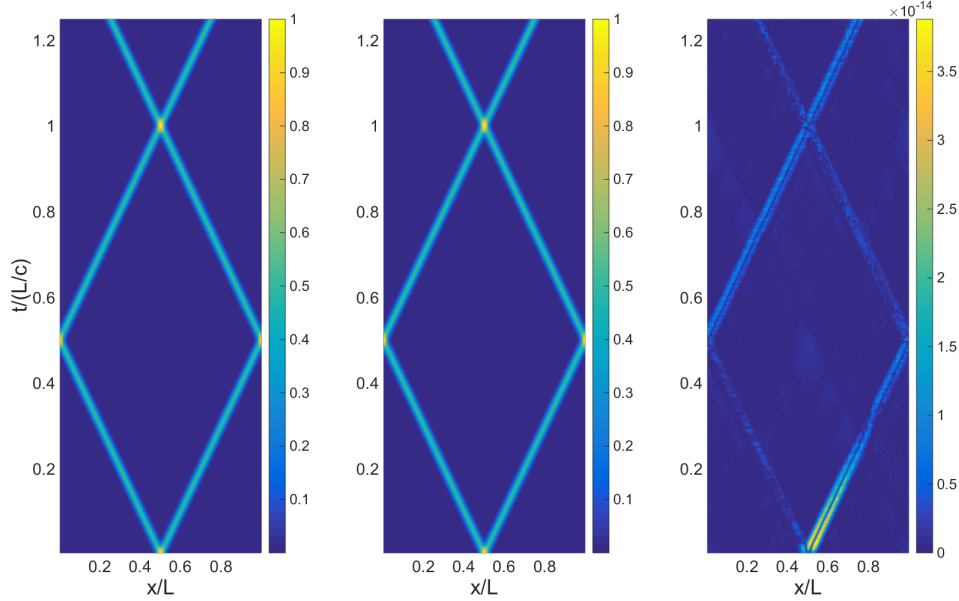


Figure 5: Two transported quantities (only density): comparison of full-order solution with sPOD approximation (left: full-order solution, middle: sPOD approximation (2 modes), right: error).

approximation matches the full solution excellently. This is also reflected in Fig. 7 where the mean error is plotted over the number of POD modes. To obtain the same accuracy as the two sPOD modes (error less than 10^{-13} , cf. Fig. 4), more than 80 POD modes are required. Furthermore, if only two POD modes are used, the relative mean error is almost 1. Of course, this insufficient performance of the POD was to be expected, since deliberately we have chosen an example, which provides a big challenge for the POD. Nevertheless, this simple example gives a first impression of the potential of the proposed sPOD.

Up to now, the reduction was based on the density alone. When considering the density and the velocity together, the convergence of the sPOD algorithm appears to be faster even though the eventually achieved error is higher (see Fig. 4). The first observation may be related to the additional available information including the coupling of density and velocity. This additional information probably leads to a smaller number of iterations needed to separate these structures. On the other hand, the asymptotic error is higher, which may be due to the higher amount of data and the related increased round-off errors occurring during the numerical computations. Apart from the convergence behavior, the results for the case that density and velocity are considered together are very similar to the results shown before and omitted therefore.

2.3 Determination of the shift velocities

It is instructive to examine the singular value spectrum of the shifted snapshot matrix as a function of the shift velocity, i.e., the (constant) velocity the time-dependent shift is based on. In Fig. 8, left, the spectrum of shifted snapshot matrices of the example considered in section

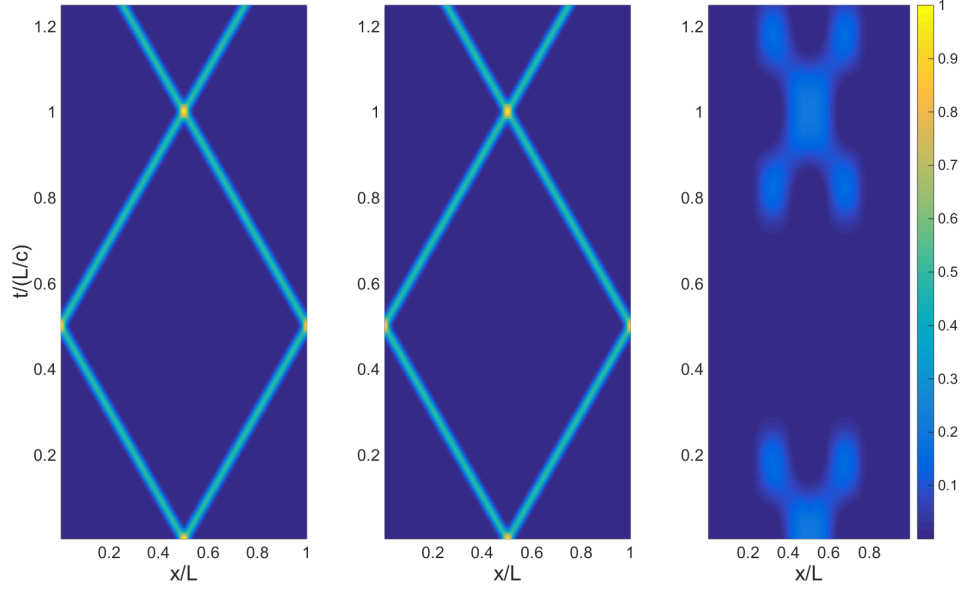


Figure 6: Two transported quantities (only density): comparison of full-order solution with sPOD and POD approximation (left: full-order solution, middle: sPOD approximation (2 modes), right: POD approximation (2 modes))

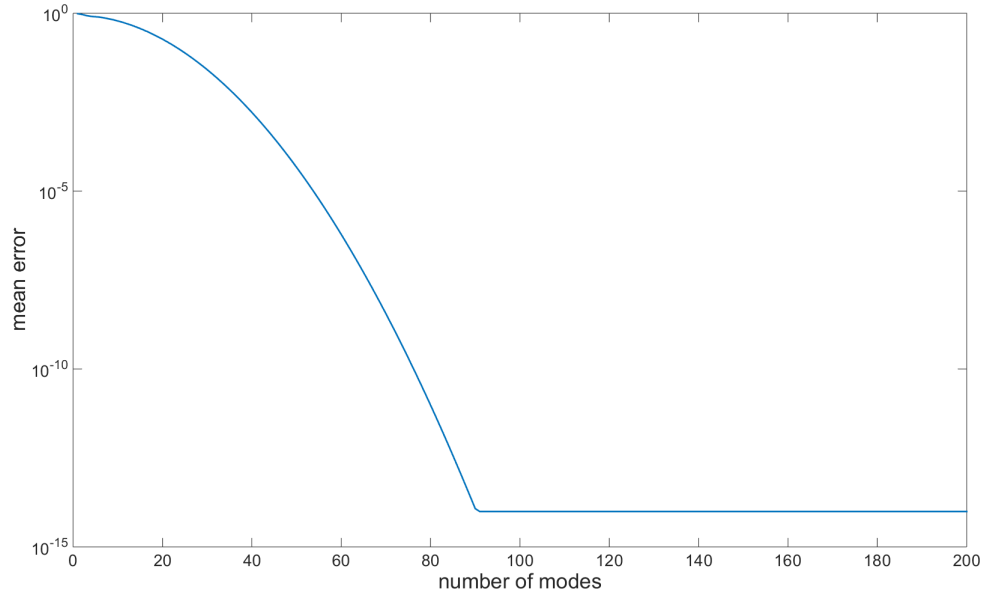


Figure 7: Two transported quantities (only density): mean error over number of POD modes

2.2 is shown for a range of shift velocities between -1.25 and 1.25 . First, it should be noted that the square integral of the solution

$$\int_0^{t_{end}} \int_0^L (q(x, t))^2 dx dt$$

does not change by a shift in x -direction. A numerical shift of the snapshot matrix should keep this invariance approximately. Thus, the Frobenius norm, which is directly connected with the singular values by

$$\|X\|_F^2 = \sum_{i,j} X_{i,j}^2 = \sum_i \sigma_i^2,$$

is conserved up to the interpolation error. This allows to directly compare spectra for different shift velocities. If a certain singular value increases by a change of shift velocity, others have to decrease to keep the sum of the squared singular values constant. The transport velocities, which are the positive and the negative sound velocity $c_{\pm} = \pm 1$, are obvious from the maxima of the leading singular value. The classical POD is recovered for a shift velocity of zero, where a slow decay of the singular values is found. The shift dependence of the pressure pulse case is contrasted with the case of a standing wave, Fig. 8, right. Here the zero shift velocity leads to a maximization of the leading singular value. However, the sound velocities are also visible as local maxima. The standing wave can be represented by two traveling waves with velocities $\pm c$, so that these are reasonable candidates for shift velocities. It should be noted, that the spectra are only a first indication. If multiple transport velocities are involved, each of the transported quantities influences the singular spectra of the others. E.g. the pressure pulse can be expressed by one mode per transport velocity, which is not obvious from Fig. 8, left, since each of the transported quantities slows the decay of the singular spectrum of the other one. The idea of shifting the snapshot matrix to obtain a stronger decay of its singular values is also used in the next subsection.

3 Two Dimensions

In section 2 we have considered a simple one-dimensional test example, which nevertheless provides a big challenge for the classical POD. The sPOD performs excellently on this example, since its structure is ideal for the sPOD algorithm. This is because the solution can exactly be represented by two moving modes with constant velocities.

In this section we want to explore the capabilities of the introduced sPOD algorithm by considering a two-dimensional transport-dominated example with non-trivial velocities. We consider a flow governed by the incompressible Navier-Stokes-equations with a vortex pair as an initial condition. The initial conditions are created from two single vortices of the form

$$\omega_{0,i} = \omega_{e,i} (1 - (r_i/r_0)^2) \exp(-(r_i/r_0)^2)$$

where $i = 1, 2$, $r_i = \sqrt{(x - x_{0,i})^2 + (y - y_{0,i})^2}$ is the distance from the respective vortex core and r_0 denotes the vortex size. The size of each vortex is chosen to be $r_0 = 0.1$ and the centers are at $(x_{0,1/2}, y_{0,1/2}) = (\pm 0.1, 0)$. The strengths are $\omega_{e,1/2} = \mp 299.5$ and the viscosity of the fluid is $\nu = 1/Re = 1/1000$. We use a *periodic* domain to avoid the further complication of the boundary treatment for the sPOD. The application of the sPOD to systems with more

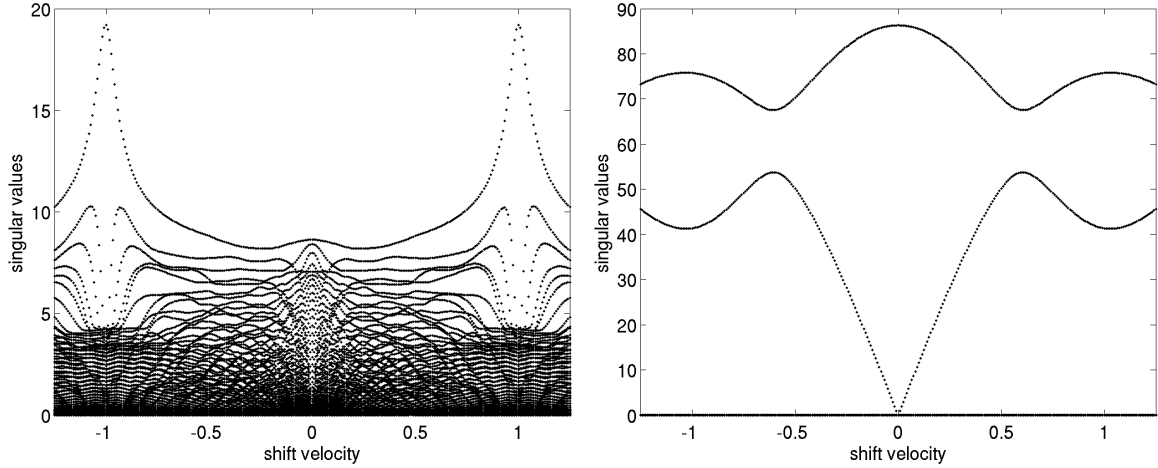


Figure 8: The singular value spectrum of the density snapshot matrix as a function of the shift velocity. Left: The pressure pulse, which creates two traveling waves with the speed $c_{\pm} = \pm 1$, which are clearly visible as maxima of the leading singular value. Right: The standing wave. The zero velocity has the biggest leading singular value. The positive and the negative sound speed are also visible as local maxima.

realistic boundaries and boundary conditions is currently under investigation. The dynamical behavior of the vortex pair is simulated by means of the energy conserving, skew-symmetric scheme described in [22], however the kind of solver is expected to have very little impact on the results of the sPOD. The periodic domain is discretized by 512^2 equidistant points with a time step of $\Delta t = 8 \cdot 10^{-4}$ s with a fractional-step time stepper. The solution is depicted in Fig. 9 by means of a contour plot for different times. One can see that the initial vorticity field induces a movement of the vortex pair in positive y -direction. Additionally, a secondary, weaker vortex pair goes in negative y -direction with a smaller transport velocity. This secondary vortex pair has motivated to choose this test case, since thereby two different, non-trivial velocities are present. The boundaries are again periodic and the simulated time is chosen such, that the primary and the secondary vortex pair do not meet (again). This example provides a significantly larger challenge for the sPOD for several reasons. First, we do not know the analytic solution of this test case. Second, the transport velocities are non-constant and unknown a priori. Third, this is a two-dimensional problem, so the amount of data to be handled is much higher than in the one-dimensional test case presented in section 2.

Despite all these new challenges we do not change the sPOD procedure but instead we perform the same algorithm as presented in section 2. Hence, we are looking for two dominant transport velocities which we assume to be constant, even though this is not strictly true. To obtain proper candidates for the transport velocities, we again consider the singular values of shifted snapshot matrices where we only shift in y -direction. Shifting in x - and y -direction should identify the y -direction as dominant, so that this extra information restriction just helps to reduce the computational time. Since the considered case is two-dimensional, we need to reshape the snapshot matrix but this is not cost-intensive in comparison to the sPOD algorithm, especially to step 6, which dominates the computation time. By a maximization of the leading singular value, we obtain $v_1 = 3$ and $v_2 = -1$ as good candidates for the dominant

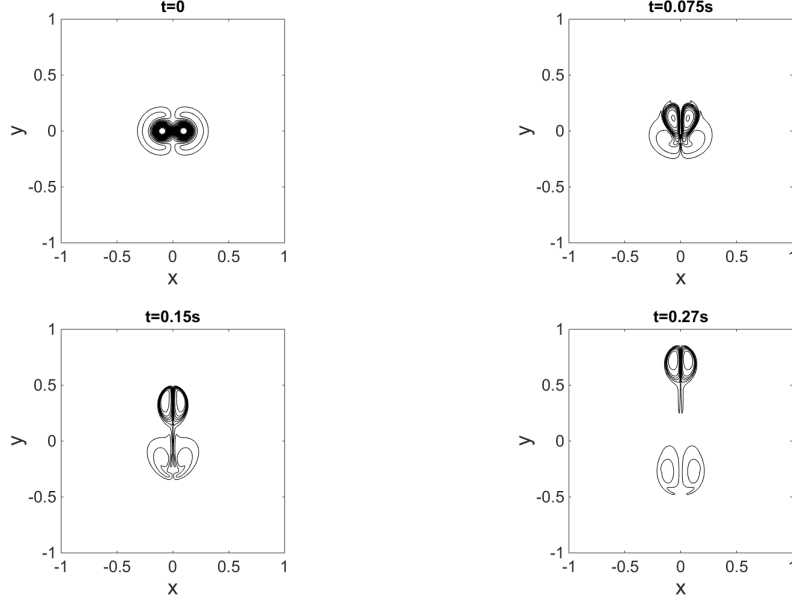


Figure 9: 2D vortex pair: full order solution (vorticity contour lines at $-270, -250, \dots, 270$)

transport velocities of the primary and secondary vortex pair, respectively. Values very close to these are obtained, if the maximization is done for a time series after the separation of the vortex pairs. Considering all snapshots results in a lower velocity and a worse reduction for the same number of modes.

A decent approximation could be obtained by choosing 10 modes per transported vortex pair, so in total 20 modes. The convergence of the sPOD for the two-dimensional vortex pair problem is depicted in Fig. 10. In comparison to the one-dimensional test example (cf. Fig. 4), the convergence is worse but still the relative mean error decreases to a value of slightly more than 1%. Due to the assumed constant velocity and the higher complexity of the transport phenomenon, a much smaller error was not to be expected.

The first three modes obtained by the sPOD for each transported vortex pair are shown in Fig. 11. The depicted modes correspond to the respective co-moving frames, i.e., $[U_\Phi]_i$ (cf. step 7 of the sPOD algorithm in section 2.2). For both velocity frames the respective first mode obviously describes the corresponding vortex pair. The higher modes seem to fulfill different roles in both cases. When considering the dynamical numerical solution, we observe that the secondary, down-going pair strongly changes in time; a rotating vorticity distribution is visible at creation. It looks very much like the second mode, thus, the higher modes represent the change of shape of the vortex pair. The shape of the primary, up-going vortex pair changes less in time. Right after the separation the sharp vortex structure emerges. The higher modes of the primary vortex pair have a strong weight at the front and rear edge of the pair. These structures seem to account for a velocity correction, since they induce a translational shift of the primary vortex pair. The velocity correction is caused by the non-constant propagation velocity of the vortex pair.

The corresponding time amplitudes of the spatial modes are depicted in Fig. 12. The amplitudes are given by the right singular vectors $[V_\Phi]_i$, describing the temporal development

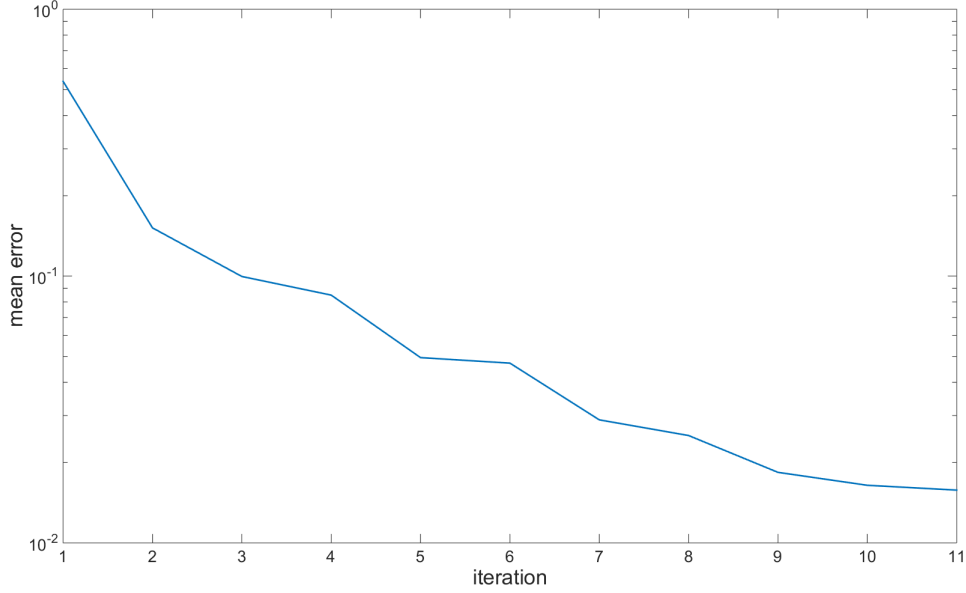


Figure 10: 2D vortex pair: convergence of the sPOD

of the respective mode $[U_\Phi]_i$, multiplied by the corresponding singular value $\sigma_{\Phi,i}$ (cf. step 7 of the sPOD algorithm in section 2.2). For large values of t , the first mode dominates the behavior of the primary vortex pair, while for the secondary pair the first two modes have a nearly similar weight. Thus, the first and the second mode of the secondary vortex pair both contribute to its shape. This can be comprehended by considering Fig. 11, where it can be seen that the vortex pair appearing in the first mode seems to fit into the cavity formed by the structures of the second mode. This agrees with the observation that the secondary vortex pair depicted in Fig. 9 is wider than the vortex pair shown in the first mode (cf. Fig. 11, bottom left). The separation of the two vortex pairs is completed at about $t > 0.15$ s, see Fig. 9. The interplay before is complicated, but afterwards a simpler picture emerges. For the primary vortex pair, the absolute value of the amplitude of the dominating first mode decays gradually in time, which reflects the reduction of the vortex strength by viscosity. The influence of the second mode becomes stronger from roughly $t = 0.2$, i.e., after the full separation. It leads to a reduction of the vortex pair transport velocity. Also, the third mode contributes to this velocity decrease. In contrast, the first time amplitude of the secondary vortex pair decreases and changes the sign over time. After separation all three amplitudes are nearly constant; this corresponds to the plausible observation that the friction has less impact on the secondary than on the primary vortex pair due to the lower velocity gradient.

While the modes give a clear and intuitive description of the flow dynamics, still some unexpected structures are visible. Namely, one can also see some stripe structures in y -direction, especially at the first mode of the secondary vortex pair (cf. Fig. 11). Our first analysis indicates that this is *not* due to an incomplete separation, but rather caused by a non-uniqueness of the decomposition. This non-uniqueness can be comprehended by the following consideration. If we add a constant offset to one of the transported quantities, the combined sPOD approximation remains unchanged if the negative counterpart offset is added to the

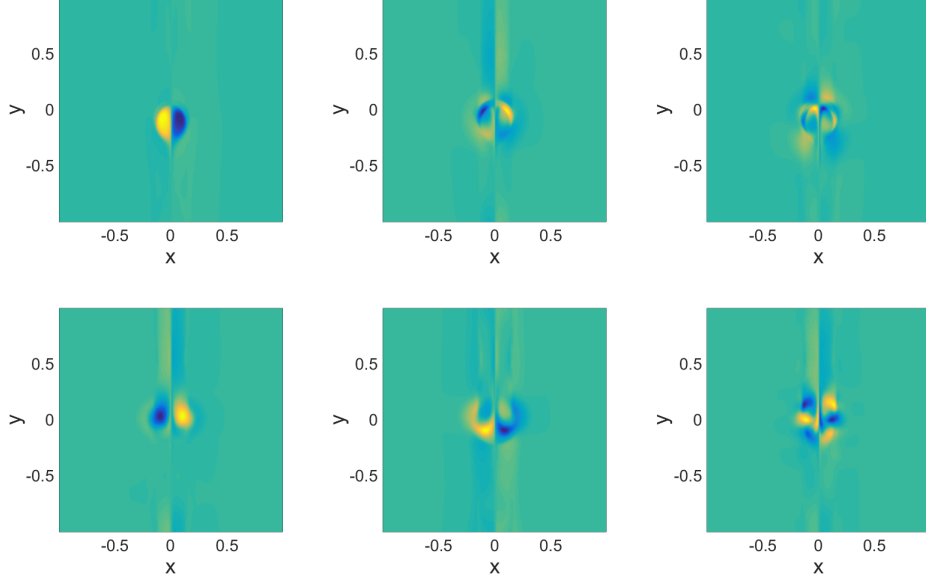


Figure 11: 2D vortex pair: first three sPOD modes for each transported vortex pair (top: primary vortex pair, bottom: secondary vortex pair)

other transported quantity. Note that offsets are shift-invariant and, hence, well represented in all velocity reference frames. The occurring stripes seem to be localized variants of these offsets along the y -direction. These unphysical structures may be removed by an additional, physically motivated constraint. This is currently being investigated. Nevertheless, the main structure of the vortex pairs is already captured quite well by the respective first modes.

The spatial modes and their time amplitudes both provide physically meaningful insights into the dynamics of the vortex movements. Finally, the sPOD is to be evaluated quantitatively by comparing it to the full order solution and the POD approximation. To this end, Fig. 13 shows the full order solution, the sPOD approximation (2×10 modes), and the POD approximation (20 modes), each for a constant time $t = 0.27s$. It is hard to see a difference between the sPOD approximation and the full order solution, since they match almost exactly. In contrast, even though the POD approximation captures the main structures of the vortex pairs quite well, it reveals some unphysical structures in the inside of the primary vortex pair (top) and, furthermore, some spurious edges in the secondary vortex pair (bottom). The superiority of the sPOD algorithm can also be comprehended by considering Fig. 14 where the mean error is plotted over the number of modes for the POD and the sPOD. One can see that the sPOD mean error is significantly smaller than the POD mean error for all numbers of modes although the difference is not as striking as in the one-dimensional test example (cf. section 2.2). This is not surprising, since the 1D example is ideal for the reduction by the sPOD while the vortex pair problem provides some additional challenges as explained in the beginning of this section. Further, the small transport distance compared with the relatively large interaction time at the beginning reduces the difference between the methods.

All in all, the sPOD method performs very well despite the non-trivial and non-constant velocities of the two-dimensional vortex pair problem and further delivers modes which allow

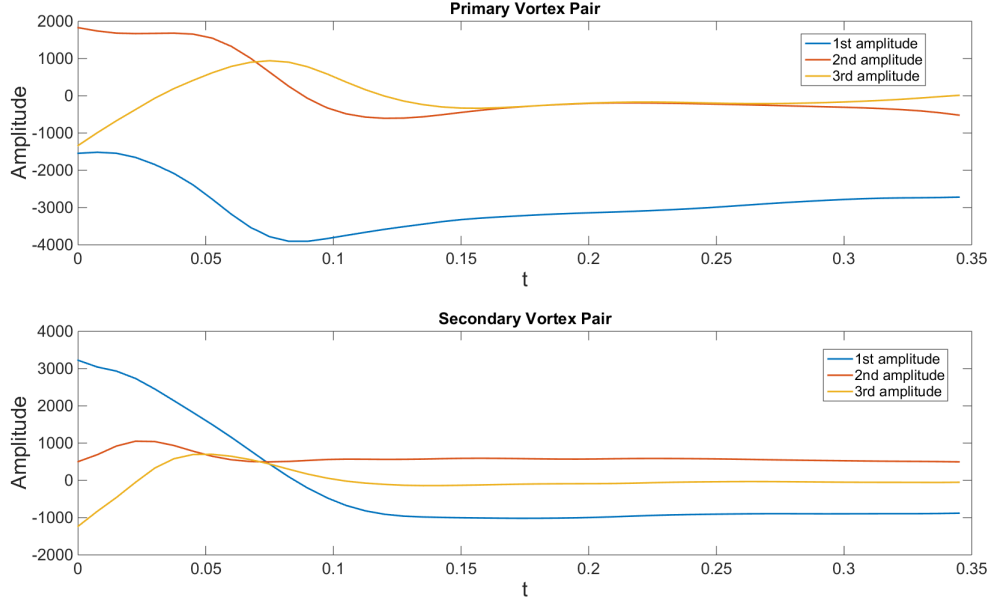


Figure 12: 2D vortex pair: time amplitudes for the first three modes for each transported vortex pair.

a physically meaningful interpretation. However, there are still a lot of aspects about the sPOD which need to be improved. These are discussed in section 4.

4 Summary and outlook

We have presented a new model reduction method for transport-dominated phenomena. The proposed method, sPOD, is a generalization of the common POD approach and, hence, delivers good low-dimensional representations for problems where the POD is applied successfully. To improve the quality of low-dimensional representations for systems where convective transport is dominant a time-dependent shift is introduced in order to obtain a co-moving variant of the snapshot matrix. Especially, a procedure is included which allows for the separation of multiple transports within the system. Furthermore, the dominant transport velocities of the system are determined by maximizing the leading singular value over different shift velocities.

We started considering the performance of the new method by means of a simple one-dimensional moving pressure pulse. Even though it is a simple phenomenon, whose analytic solution can be written as two moving modes with constant velocities, the common POD approach is not able to yield a low-dimensional representation of the solution, which captures the dynamical behavior reasonably well. In contrast, the sPOD performs excellently representing the solution with only two modes, which capture the dynamics almost perfectly.

After this simple one-dimensional test example, we considered a more challenging case of a two-dimensional moving vortex pair. The occurring transport velocities of the primary and the secondary vortex pair are non-trivial and non-constant. Despite of these difficulties, the sPOD algorithm provides a physically meaningful low-dimensional representation which captures the dynamics excellently. Moreover, it outperforms the standard POD which yields

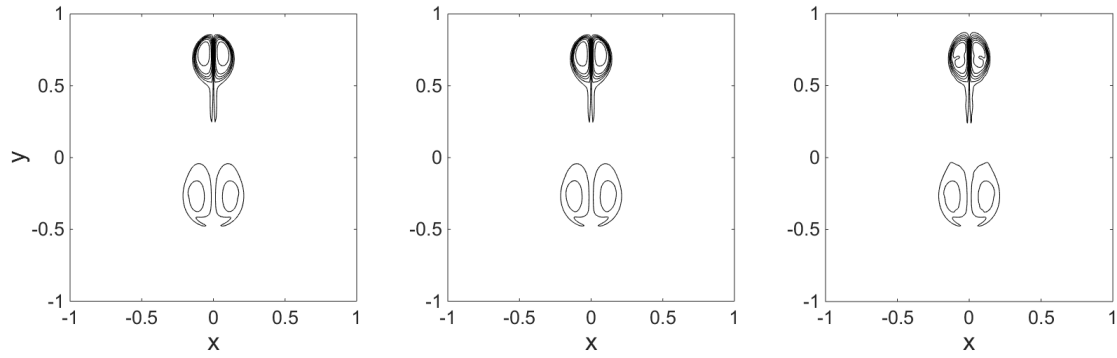


Figure 13: 2D vortex pair: comparison of full-order solution with sPOD and POD approximation (left: full-order solution, middle: sPOD approximation (20 modes), right: POD approximation (20 modes), $t = 0.27s$, vorticity contour lines at $-270, -250, \dots, 270$)

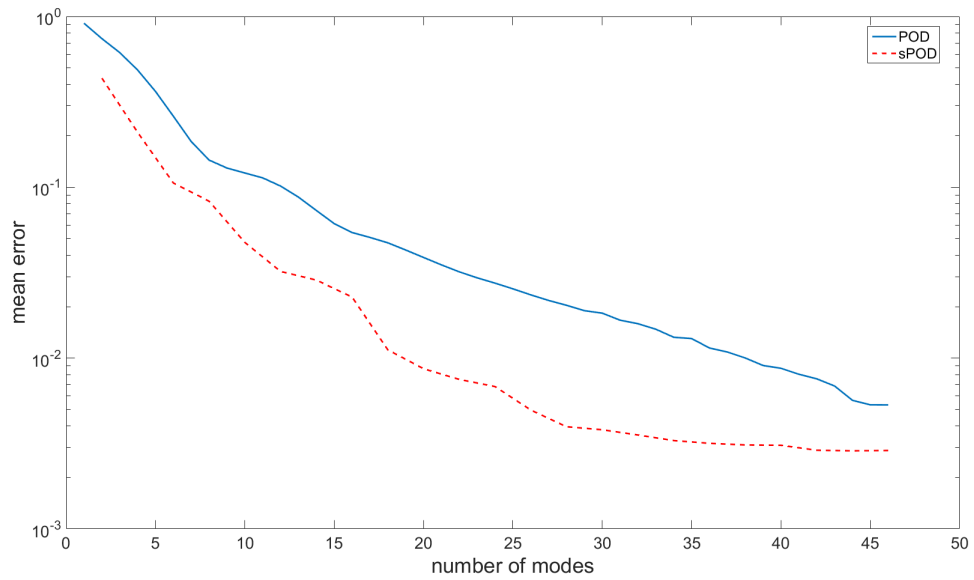


Figure 14: 2D vortex pair: mean error over number of modes (POD vs sPOD)

a low-dimensional representation suffering from unphysical structures and spurious edges.

The results and observations presented in this paper motivate for further development of the proposed sPOD algorithm. First, the algorithm itself is still rather heuristic and needs mathematically more profound understanding, especially regarding aspects like convergence behavior and error estimation. In addition, up to now we have only considered systems with periodic boundaries. However, in practical applications, one often faces systems with boundaries and corresponding boundary conditions. How to account for these in the shifting procedure will be subject of future research to extent the applicability of the sPOD to systems with physically meaningful boundary conditions. Another aspect that has to be considered in more detail is the occurrence of non-constant velocities. The considered vortex pair already exhibits this feature but explicitly accounting for a temporal change in the transport velocities might be a promising extension of the sPOD to improve its performance when dealing with non-constant transport velocities. In this context, the constant velocity assumption used for the shifting procedure up to now needs to be dropped. Further, while a good separation was found, a non-intuitive stripe structure occurred in the modes. Its origin and corresponding improvements are currently under investigation. Finally, the efforts presented up to now just focus on finding a low-dimensional representation of the full order solution. To use the corresponding ansatz spaces for reduced order models either by interpolation or by Galerkin projection of the full order model is to be investigated as well. This leads to reduced-order models, which can well be applied for control, optimization, and design.

Acknowledgements. The authors gratefully acknowledge the support by the Deutsche Forschungsgemeinschaft (DFG) as part of the collaborative research center SFB 1029 *Substantial efficiency increase in gas turbines through direct use of coupled unsteady combustion and flow dynamics*, project A02 *Development of a reduced order model of pulsed detonation combustor*.

References

- [1] A. C. Antoulas, C. A. Beattie, and S. Gugercin. Interpolatory model reduction of large-scale dynamical systems. In J. Mohammadpour and K. M. Grigoriadis, editors, *Efficient modeling and control of large-scale systems*, pages 3–58. Springer New York, USA, 2010.
- [2] T. J. Barber, M. H. Ahmed, and N. Y. Abdel Shafi. POD snapshot data reduction for periodic fluid flows. In *43rd AIAA Aerospace Sciences Meeting and Exhibit*, 2005.
- [3] C. Beattie and S. Gugercin. Interpolatory projection methods for structure-preserving model reduction. *Systems & Control Letters*, 58(3):225–232, 2009.
- [4] W. J. Beyn and V. Thümmler. Freezing solutions of equivariant evolution equations. *SIAM Journal on Applied Dynamical Systems*, 3(2):85–116, 2004.
- [5] K. K. Chen, J. H. Tu, and C. W. Rowley. Variants of dynamic mode decomposition: Boundary condition, Koopman, and Fourier analyses. *Journal of Nonlinear Science*, 22(6):887–915, 2012.

- [6] P. G. Constantine and G. Iaccarino. Reduced order models for parameterized hyperbolic conservation laws with shock reconstruction. Annual research briefs, Center for Turbulence Research, Stanford, USA, 2012.
- [7] P. Deheuvels and G. V. Martynov. A Karhunen–Loeve decomposition of a Gaussian process generated by independent pairs of exponential random variables. *Journal of Functional Analysis*, 255(9):2363–2394, 2008.
- [8] M. A. Dihlmann and B. Haasdonk. Certified PDE-constrained parameter optimization using reduced basis surrogate models for evolution problems. *Computational Optimization and Applications*, 60(3):753–787, 2015.
- [9] J. L. Eftang. *Reduced basis methods for parametrized partial differential equations*. PhD thesis, Norwegian University of Science and Technology, Trondheim, Norway, 2011.
- [10] F. Fedele, O. Abessi, and P. Roberts. Symmetry reduction of turbulent pipe flows. *ArXiv e-prints*, 2015. <http://arxiv.org/abs/1412.6711v3>.
- [11] K. Fujimoto and J. M. A. Scherpen. Balanced realization and model order reduction for nonlinear systems based on singular value analysis. *SIAM Journal on Control and Optimization*, 48(7):4591–4623, 2010.
- [12] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, USA, fourth edition, 2013.
- [13] M. A. Grepl. Certified reduced basis methods for nonaffine linear time-varying and nonlinear parabolic partial differential equations. *Mathematical Models and Methods in Applied Sciences*, 22(03):1150015, 2012.
- [14] S. Gugercin, A. C. Antoulas, and C. A. Beattie. A rational Krylov algorithm for optimal H2 model reduction. In *Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems*, pages 1665–1667, 2006.
- [15] M. Hinze and S. Volkwein. Proper orthogonal decomposition surrogate models for nonlinear dynamical systems: Error estimates and suboptimal control. In P. Benner, D. C. Sorensen, and V. Mehrmann, editors, *Dimension reduction of large-scale systems*, volume 45 of *Lecture Notes in Computational Science and Engineering*, pages 261–306. Springer Berlin Heidelberg, Germany, 2005.
- [16] I. T. Jolliffe. *Principal Component Analysis*. Springer New York, USA, 1986.
- [17] M. Lemke, A. Miedlar, J. Reiss, V. Mehrmann, and J. Sesterhenn. Model reduction of reactive processes. In R. King, editor, *Active Flow and Combustion Control 2014*, volume 127 of NNFM, pages 234–262. Springer International Publishing, Cham, Switzerland, 2015.
- [18] D. M. Luchtenburg, B. R. Noack, and M. Schlegel. An introduction to the POD Galerkin method for fluid flows with analytical examples and MATLAB source codes. Technical Report 01/2009, Department of Fluid Dynamics and Engineering Acoustics, TU Berlin, 2009.

- [19] D. J. Lucia. *Reduced order modeling for high speed flows with moving shocks*. PhD thesis, Air Force Institute of Technology, Ohio, USA, 2001.
- [20] A. J. Mayo and A. C. Antoulas. A framework for the solution of the generalized realization problem. *Linear Algebra and its Applications*, 425(2-3):634–662, 2007.
- [21] M. Rathinam and L. R. Petzold. A new look at proper orthogonal decomposition. *SIAM Journal on Numerical Analysis*, 41(5):1893–1925, 2003.
- [22] J. Reiss. A family of energy stable, skew-symmetric finite difference schemes on collocated grids. *Journal of Scientific Computing*, 65(2):821–838, 2015.
- [23] C. W. Rowley, I. G. Kevrekidis, J. E. Marsden, and K. Lust. Reduction and reconstruction for self-similar dynamical systems. *Nonlinearity*, 16(4):1257–1275, 2003.
- [24] G. Rozza, D. B. P. Huynh, and A. T. Patera. Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations. *Archives of Computational Methods in Engineering*, 15(3):229–275, 2008.
- [25] P. J. Schmid and J. L. Sesterhenn. Dynamic mode decomposition of numerical and experimental data. In *61st Annual Meeting of the APS Division of Fluid Dynamics*, 2008.
- [26] J. H. Tu. *Dynamic mode decomposition: Theory and applications*. PhD thesis, Princeton University, USA, 2013.
- [27] J. H. Tu and C. W. Rowley. An improved algorithm for balanced POD through an analytic treatment of impulse response tails. *Journal of Computational Physics*, 231(16):5317–5333, 2012.
- [28] K. Urban and A. T. Patera. A new error bound for reduced basis approximation of parabolic partial differential equations. *Comptes Rendus Mathematique*, 350(3–4):203–207, 2012.
- [29] K. Willcox and J. Peraire. Balanced model reduction via the proper orthogonal decomposition. *AIAA Journal*, 40(11):2323–2330, 2002.